



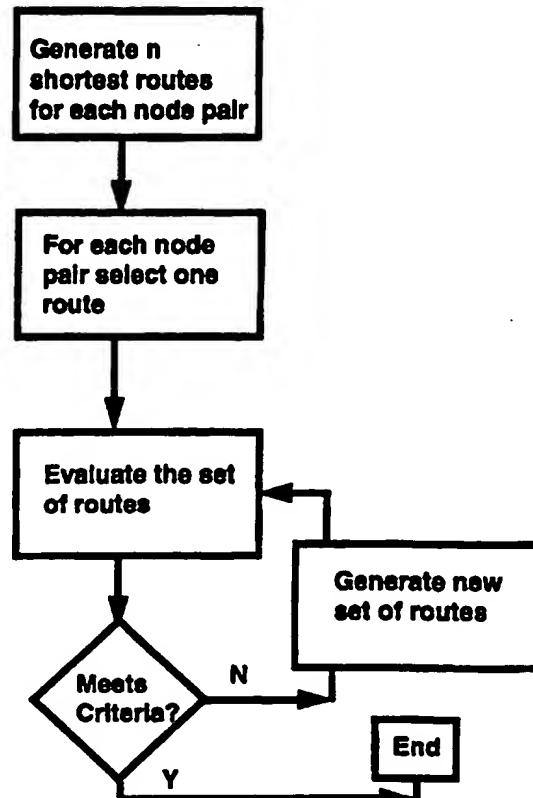
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H04L 12/56		A1	(11) International Publication Number: WO 96/31969
			(43) International Publication Date: 10 October 1996 (10.10.96)
<p>(21) International Application Number: PCT/GB96/00769</p> <p>(22) International Filing Date: 29 March 1996 (29.03.96)</p> <p>(30) Priority Data: 9506772.4 1 April 1995 (01.04.95) GB</p> <p>(71) Applicant (for all designated States except US): NORTHERN TELECOM LIMITED [CA/CA]; World Trade Center of Montreal, 8th floor, 380 St. Antoine Street West, Montreal, Quebec H2Y 3Y4 (CA).</p> <p>(72) Inventors; and</p> <p>(75) Inventors/Applicants (for US only): TURNER, John [GB/GB]; Home Cottage, Thriplow Road, Foulmere, Cambridgeshire SG8 7QT (GB). SMITH, George [GB/GB]; (GB). MANN, Jason [GB/GB]; School of Information Systems, University of East Anglia, Norwich NR4 7TJ (GB).</p> <p>(74) Agent: RYAN, John, Peter, William; Nortel Limited, Patents and Licensing, West Road, Harlow, Essex CM20 2SH (GB).</p>			(81) Designated States: US, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i>

(54) Title: TRAFFIC ROUTING IN A TELECOMMUNICATIONS NETWORK

(57) Abstract

A telecommunications network comprising a plurality of nodes interconnected via communications paths is configured by the determination of a routing plan for the traffic carried by the network. A set of possible paths, e.g. the eight shortest paths between each pair of network nodes, is determined, and each path is allocated a digital code characteristic of that path identity. Groups of paths are selected, each group comprising one path from each set, and for each said group being characterised by a bit string comprising the digital codes of the paths comprising that group. These bit strings are then used as starting values or chromosomes in a genetic algorithm search procedure to generate and evaluate new path groups so as to determine an optimum combination or routing plan of paths between the network nodes.



Best Available Copy

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroun	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

TRAFFIC ROUTING IN A TELECOMMUNICATIONS NETWORK

This invention relates to the routing of telecommunications traffic in a network comprising a plurality of nodes interconnected by paths. In particular, the invention relates to a method of defining a traffic routing plan for such a network and to a method of routing traffic via the routing plan.

5

BACKGROUND OF THE INVENTION

As telecommunications networks become more complex and the traffic load increases there is a need to determine the routing of traffic within a network so as to minimise the number of communication channels used, to 10 reduce the risk of being unable to handle traffic and to minimise the total cost to the system operator. At present, traffic routing is determined largely on an empirical basis as a full analysis of a complex network to obtain an analytical solution is an intractable problem. It is thus very difficult if not impossible for an operator to ensure that a network is used 15 with a high degree of efficiency under all traffic conditions. It is also difficult for an operator to provide efficient allocation of new traffic routes on demand.

The object of the present invention is to minimise or to overcome this 20 disadvantage.

SUMMARY OF THE INVENTION

According to the invention there is provided a method of configuring a telecommunications network whereby to determine traffic routing, said 25 network comprising a plurality of nodes interconnected by communications paths, and each said path incorporating one or more communications channels, wherein the method includes determining a set of possible paths between each pair of network nodes, allocating to each path of each said set a digital code characteristic of that path identity, selecting groups of 30 paths each said group comprising one path from each set, providing for each said group a respective bit string comprising the digital codes of the paths comprising that group, and using the bit strings of said groups as starting values or chromosomes in a genetic algorithm search procedure

5 whereby to generate new bit strings corresponding to new path groups, determining for each said new path group a fitness function representative of the effectiveness of that new path group relative to the other path groups, and evaluating the new path groups from their respective fitness functions so as to determine an optimum combination of paths between the network nodes.

10 In another aspect, the invention provides a method of routing traffic in a telecommunications network comprising a plurality of nodes interconnected by communications paths, each said path incorporating one or more communications channels, wherein the method includes determining a set of possible paths between each pair of network nodes, allocating to each path of each said set a digital code characteristic of that path identity, selecting groups of paths each said group comprising one 15 path from each set, providing for each said group a respective bit string comprising the digital codes of the paths comprising that group, and using the bit strings of said groups as starting values or chromosomes in a genetic algorithm search procedure whereby to generate and evaluate new path groups so as to determine an optimum combination of paths 20 between the network nodes, and routing the traffic via said optimum combination of paths.

25 The genetic algorithm uses the bit strings representing the paths as chromosomes to generate new combinations of paths which are then evaluated to determine those which represent better solutions to the routing problem. The bit strings of these new paths then form new starting values in an iterative procedure which may be continued either until a solution satisfies predetermined criteria or until a predetermined number of iterative cycles have been completed.

30 30 The technique may be employed in the construction of a transmission network manager to provide control of traffic routing. The technique may also be employed in ATM switching applications to provide predictive capacity management for the allocation of virtual channels for voice traffic.

35 35 A general description of the genetic algorithm procedure is given by J. H. Holland in 'Adaptation in Natural and Artificial Systems', University of Michigan Press, Ann Arbor, MI.

SIMPLE DESCRIPTION OF THE DRAWINGS

An embodiment of the invention will now be described with reference to the accompanying drawings in which;

5 Figure 1 is a schematic diagram of a typical network illustrating the routing of traffic between nodes;

Figure 2 is a flow chart illustrating the process of evaluating combinations of traffic routes;

10 Figure 3 illustrates the genetic algorithm procedure for determining routing solutions for the network of figure 1.

DESCRIPTION OF PREFERRED EMBODIMENT

15 Referring to figure 1, there is depicted a highly schematic diagram of a telecommunications network 11 comprising a number of nodes A to Q, each adjacent pair of nodes in the network being coupled via a corresponding link. Traffic between a pair of nodes, e.g. A and B, is routed over a path which comprises one or more links. For example, traffic between nodes A and B may be routed over a path comprising the links between the nodes K, L, M and J, or over a path via the nodes E, F, G, H and J. In general there will be a large number of possible paths between a pair of nodes. In the network of figure 1, each link may conveniently be identified by its end nodes. Thus, the link between the nodes K and L is represented as KL. It will be appreciated that in the network of figure 1, there may be some links whose traffic handling capacity is higher than that of other links. In the network of figure 1, the nodes may comprise switches or exchanges and the links may comprise trunk groups.

20

25

30 The choice of a particular route between a pair of nodes will in general depend on considerations of the bandwidth required to accommodate the traffic and the cost of the path to the network operator. However, this choice cannot be made in isolation as the traffic route between every other pair of system nodes must also be considered. Thus, in the network of figure 1, traffic between nodes C and D may be routed along a path via nodes N, M, L, G, P and Q. However, this particular route uses the link LM which is already used by one of the routes from A to B. If too many routes in the final routing plan for the network also use this particular link there may well be insufficient bandwidth on that link to accommodate the routing

35

plan. It will also be appreciated that in the network of figure 1, each individual link will have a designated bandwidth. In the European SDH standard format this bandwidth may be represented in the form 1x63, 4x63 or 16x63, these being the number of virtual containers that are 5 accommodated on a STM1, a STM4 or a STM16 path respectively. Thus, when considering usage of a particular network link, the bandwidth or capacity of that link must be taken into account.

It will be understood that when the routing of traffic between every pair of 10 network nodes is considered, the problem of defining an optimum routing plan becomes extremely complex. A generally preferred solution to be aimed at is one in which every system link is used to 50% of its capacity, although in practice an acceptable solution is one which approaches this preferred solution.

15 Referring now to figure 2, this comprises a flow chart illustrating our method of determining an effective routing plan for the telecommunications network of figure 1. For every pair of nodes in the network we determine a set of n routes or paths between those nodes. Preferably, these paths 20 comprise the n shortest routes between the two nodes. Conveniently n is a power of 2 and for most applications a set of 8 shortest routes may conveniently be employed. Each of these routes is identified by a three bit binary code which, in the case of 8 routes, will range from 000 to 111. The next step is to form groups of paths, each group comprising paths selected 25 one from each set. Each group is represented as a unique bit string or chain formed from the binary codes of the paths forming the group. Thus for example, the bit string 010100110..... represents the group comprising the third path between the first pair of nodes, the fifth path between the second pair of nodes, the seventh path between the third pair of nodes 30 and so on. Although the routes shown in figure 2 are point to point, it will be appreciated that the technique may be adapted to address the problem of point to multipoint routing.

35 The groups of paths thus determined represent each a possible solution to the network routing problem, although some of these possible solutions may of course be infeasible e.g. on the grounds of cost and/or of traffic handling capability. These possible solutions are evaluated to determine a measure of their effectiveness as solutions so that potentially good solutions can be retained and ineffective or poor solutions discarded. This

Evaluation is made on the basis of a fitness function whose determination will be described below.

5 The first part of the evaluation process assigns a cost, i.e. a network operator cost, to each bit string in the bit chain representing a path comprising one or more links. This cost is made up of the sum of the costs of the respective network links used by the traffic between the corresponding nodes. These costs, which may be stored in a lookup table, are summed to provide a path cost term (C) in a fitness function for the network

10

In addition to the path cost, we also consider the total utilisation of each link in a proposed path group. For each group of paths we construct a utilisation matrix or table which identifies the total traffic demand imposed

15 on every network link by those paths of the group which incorporate that link. I.e. for each path in that group we record a tally of the links used by that path. This tally will of course identify links whose capacity would be overused or underused, and for each potential solution we determine the number (B) of links whose bandwidth capability would either be exceeded or would exceed a predetermined level of usage. This is in effect a congestion factor.

20

25 The evaluation process next determines the number of channels and/or the band width that will be required to support the traffic requirements. The amount of traffic on any link determines the number of channels (N) that are required to support that traffic.

We then define a fitness function F for each potential solution as:-

$$F = w_1 \times C + w_2 \times N + w_3 \times R$$

Where w_1 , w_2 and w_3 are weighting factors which the operator can define to force the process to concentrate on minimising either the path costs or the number of channels used, or to provide a trade off between these.

35 constraints. Appropriate values of these weighting factors will be determined by the system operator to match the particular network configuration and traffic requirements. A bias against selection of particular routes can be introduced by adjustment of the corresponding weighting factor w_3 so as to reserve these routes for subsequent

availability to priority customers or so as to inhibit the routing of traffic on paths whose occupancy is already greater than a predetermined value. The purpose of the genetic algorithm procedure is to generate new solutions or path groups whose fitness function is minimised.

5

Referring now to figure 3, this illustrates the genetic algorithm process for determining routing information from the potential solutions represented by the groups of paths. Each potential solution or path group can be regarded as a chromosome and the component bit codes representing the 10 component paths of the group can be considered as the genes of that chromosome. The paths to be employed as the initial inputs to the genetic algorithm procedure may either be determined on a random basis or a preferred set of paths may be determined from their fitness function as defined above so that the most unfavourable paths are excluded from the 15 initial selection.

The function "select" in figure 3 is used to sample the population pool P to generate a gene pool Q. this selection may be done with replacement so that a good solution may be chosen many times. On the other hand, bad 20 solutions may not appear at all in Q. The desired quality in any select mechanism is that good solutions have a better than average chance of being chosen, while bad solutions have a worse than average chance of being chosen. In the described embodiment, this selection is based on the fitness function defined above.

25

If "select" were the only genetic operator available, then the population could quickly converge to multiple copies of the best solution in the initial population. In order to ensure that the search space is effectively explored, recombination operators are employed with the "create" function 30 of figure 3. Typically, we employ crossover and mutation functions described below.

Crossover

In the crossover function, solution strings are chosen at random and, 35 with a predetermined probability p_c , identically positioned substrings are exchanged. For example, in a single point crossover, a random crossing point is defined and the portions of the strings to the one side of this point are exchanged. Thus, the strings:

10101Λ10011010010

abbbaΛabbaabaabbb

are recombined to provide the new strings

10101abbaabaabbb

5 and

abbba10011010010

In a similar multipoint crossover process, a number of crossing points are selected and the corresponding bit substreams are interchanged.

10

Mutation

If we have, for example, an initial population of bit strings in which, for one particular position, the bit value is identical for all members of the population, then all offspring of this population generated by the 15 crossover process will retain this characteristic. This will obstruct exploration of the full search space. To overcome the risk that this situation might occur, we employ a mutation process to introduce random changes at the gene level, i.e. at the bit level. A gene or bit is selected at random with a probability p_m , and its value is changed from 20 ONE to ZERO or from ZERO to ONE. This introduces a background mutation into the evolutionary process and ensures that the entire search space is made available. Operation of the genetic algorithm on this mutating population ensures that undesirable mutations die out, but desirable mutations are propagated through the population in 25 subsequent generations.

Other, application specific operators may also be employed to ensure that the search space is as extensive as possible and that the better solutions are retained in the population pool.

30

The genetic algorithm (figure 3) commences with an essentially random population of chromosomes, each chromosome representing a possible solution to the routing problem. Some of these chromosomes are selected by the algorithm on the basis of their fitness function defined above to form 35 a gene pool **Q** from which a new set or offspring pool **R** of chromosomes is created using a genetic operator. A new population pool **P'** is then constructed from the original pool **P**, the gene pool **Q** and the offspring pool **R**. This process is continued either until a solution is found which satisfies a predetermined criterion determined from the fitness function or

until a particular time period has elapsed after which it is expected that little further improvement will result. It will be appreciated that the solution obtained in this manner may or may not be the best possible analytic solution to the routing problem, but it will at worst be a close approximation thereto. In formal terms, the algorithm may be listed as follows:-

```
GA
{objective is to minimise value (p) such that p ∈ U}
P, Q, R: multiset of solutions ⊂ U;
10    initialise (P)
      while not finish (P) do
          begin
              Q := select (P);
              R := create (Q);
15          P := merge (P,Q,R)
          end
      end GA
```

20 In each new iteration, the algorithm manipulates the genes in the gene pool to generate corresponding new path combinations which are then evaluated for fitness. By selecting the better potential solutions this establishes a gene pool containing the more favourable genes. The effects of the crossover and mutation functions can be varied as the algorithm proceeds by suitable adjustment of their respective probability
25 values p_c and p_m .

Having determined an appropriate solution to the problem in the form of a binary bit string, this solution is then translated back into the corresponding path routing information for application to the network.

30 In a development of the technique, the search procedure may be facilitated by "seeding" the population pool of a new search with solutions found from a previous similar search and stored in a memory. We have found that this injection of previously found routing patterns into the initial population
35 employed for a new search can reduce the number of iterations required to converge on an acceptable solution and thus reduce the time required to reach this solution.

in a further application, the technique may be employed as a design tool for the modification or extension of an existing network e.g. to accommodate increased traffic demand. Such modification will generally include the addition of further nodes to the network and/or an increase in
5 the channel capacity of one or more network links.

CLAIMS

1. A method of configuring a telecommunications network whereby to determine traffic routing, said network comprising a plurality of nodes interconnected by communications paths, and each said path incorporating one or more communications channels, wherein the method includes determining a set of possible paths between each pair of network nodes, allocating to each path of each said set a digital code characteristic of that path identity, selecting groups of paths each said group comprising one path from each set, providing for each said group a respective bit string comprising the digital codes of the paths comprising that group, and using the bit strings of said groups as starting values or chromosomes in a genetic algorithm search procedure whereby to generate new bit strings corresponding to new path groups, determining for each said new path group a fitness function representative of the effectiveness of that new path group relative to the other path groups, and evaluating the new path groups from their respective fitness functions so as to determine an optimum combination of paths between the network nodes.
2. A method as claimed in claim 1, wherein said fitness function associated with a path is determined from the total cost, the total number of communication channels used and the traffic congestion associated with that path.
3. A method as claimed in claim 2, wherein the components of the fitness function are individually weighted
4. A method as claimed in any one of claims 1 to 3, wherein each said set of paths comprises the shortest routes between the respective nodes.
5. A method as claimed in any one of claims 1 to 4, wherein said communications channels are synchronous channels
6. A method as claimed in any one of claims 1 to 4, wherein said communications channels are virtual channels in an ATM network.
7. A method as claimed in any one of claims 1 to 6, wherein the number of paths in each said set is a power of two.

8. A method as claimed in any one of claims 1 to 7, wherein said genetic algorithm provides crossover and mutation of said bit strings.
9. A method as claimed in claim 7, wherein each said set of paths
5 comprises eight paths.
9. A telecommunications network configured by a method as claimed
in any one of the preceding claims.
10. A method of routing traffic in a telecommunications network
comprising a plurality of nodes interconnected by communications paths,
each said path incorporating one or more communications channels,
wherein the method includes determining a set of possible paths between
each pair of network nodes, allocating to each path of each said set a
15 digital code characteristic of that path identity, selecting groups of paths
each said group comprising one path from each set, providing for each
said group a respective bit string comprising the digital codes of the paths
comprising that group, and using the bit strings of said groups as starting
values or chromosomes in a genetic algorithm search procedure whereby
20 to generate and evaluate new path groups so as to determine an optimum
combination of paths between the network nodes, and routing the traffic
via said optimum combination of paths.

1/3

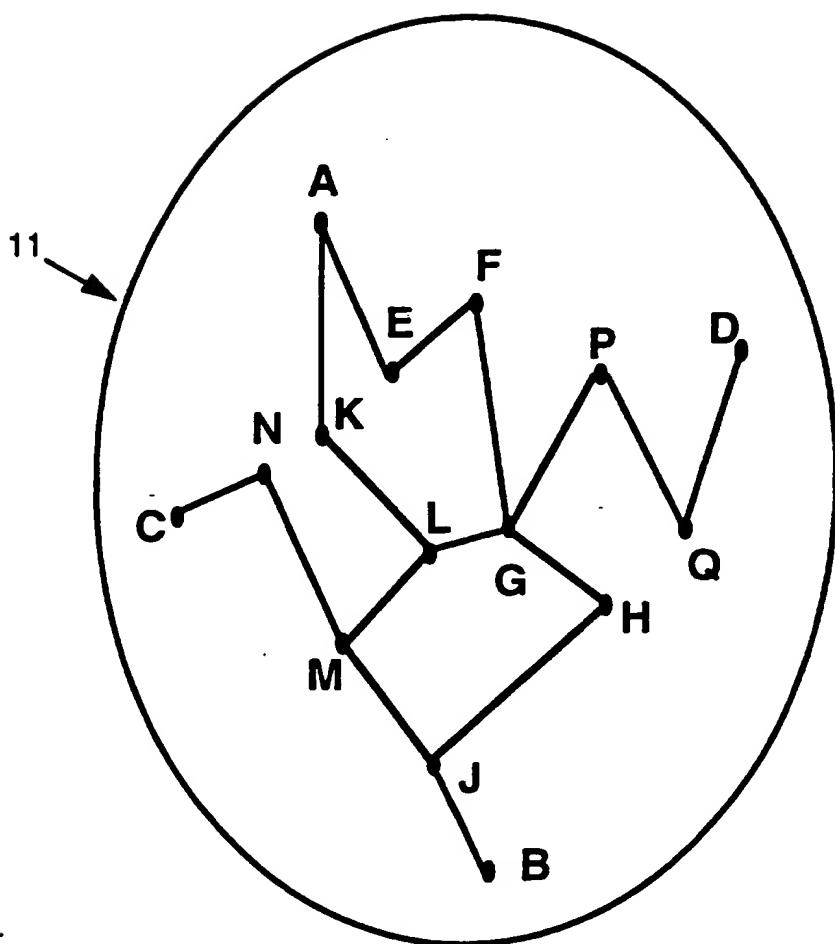
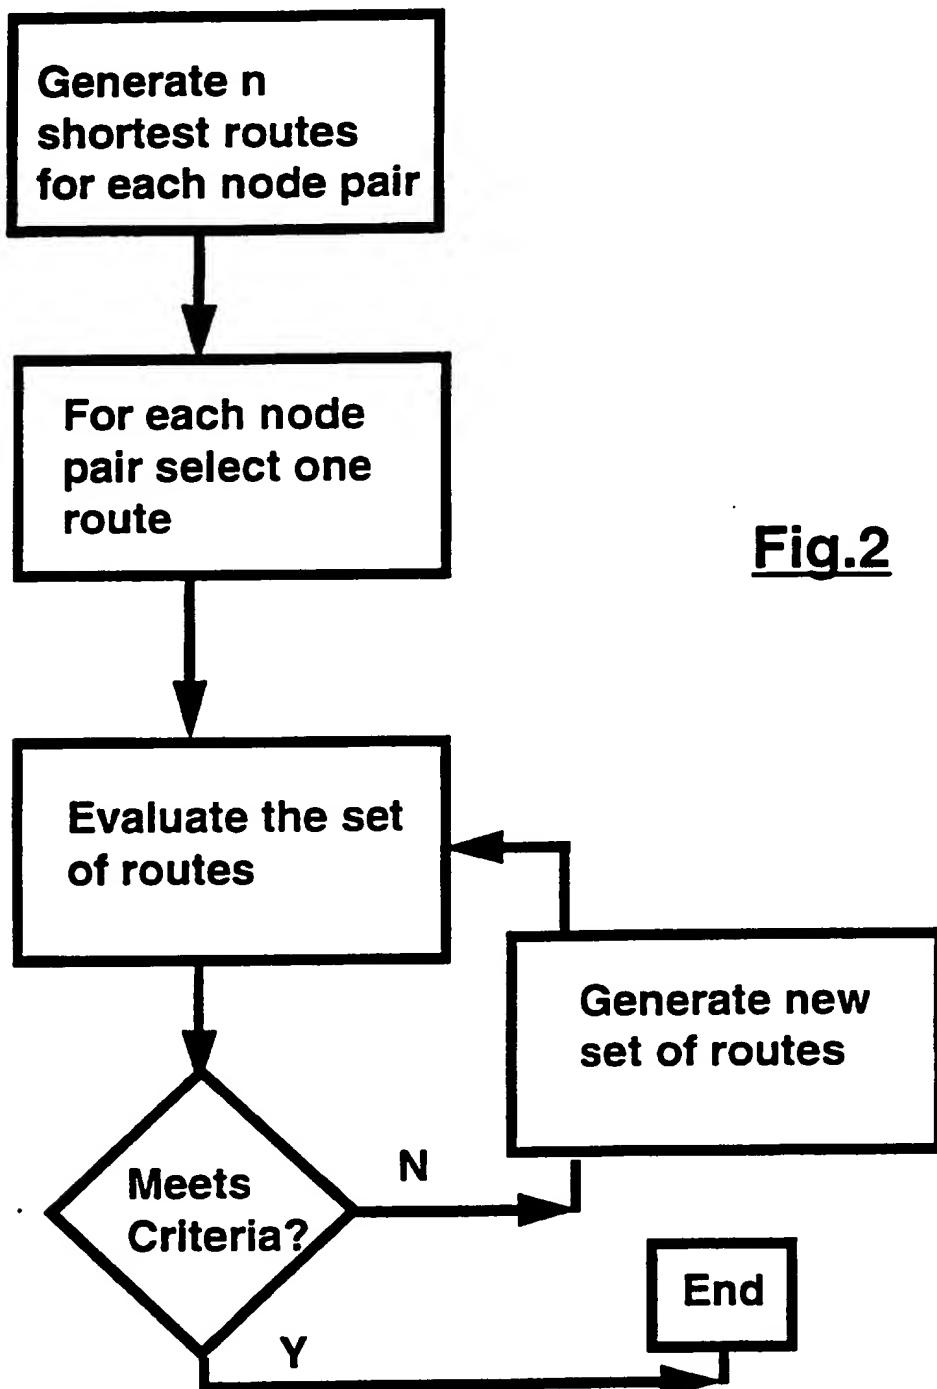


Fig.1



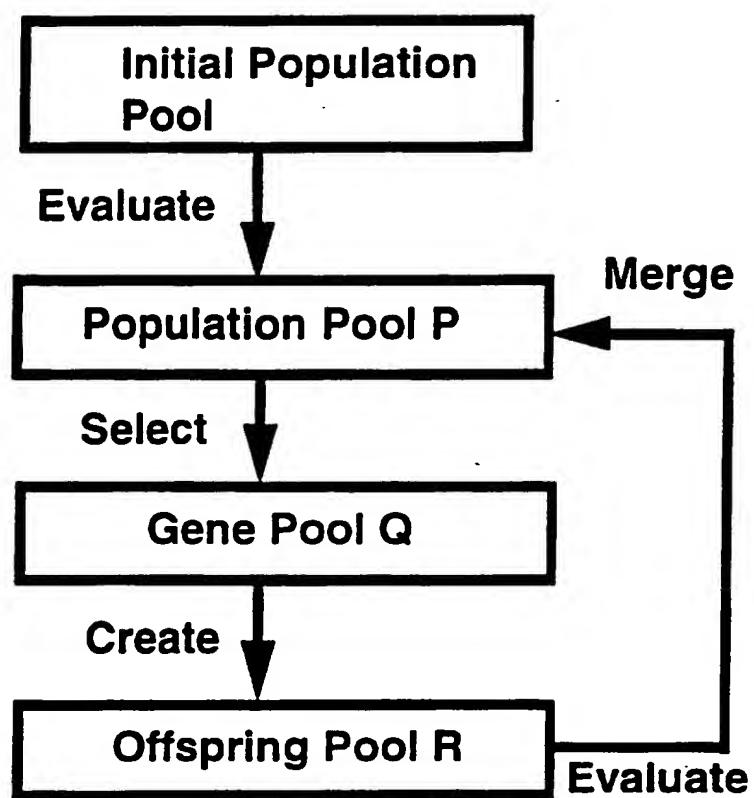


Fig.3

INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB 96/00769

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	INTERNATIONAL CONFERENCE ON NEURAL NETWORKS, vol. 2, April 1993, USA, pages 1123-1128, XP000371055 N. SHIMAMOTO ET AL.: "A dynamic routing control based on a genetic algorithm" see the whole document	1,4,6,8, 10
Y	---	2,3
A		5,7,9
Y	US,A,5 291 477 (S.C. LIEW) 1 March 1994 see claims ---	2,3 -/-

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- *'A' document defining the general state of the art which is not considered to be of particular relevance
- *'E' earlier document but published on or after the international filing date
- *'L' document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *'O' document referring to an oral disclosure, use, exhibition or other means
- *'P' document published prior to the international filing date but later than the priority date claimed

- *'T' later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *'X' document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *'Y' document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- *'A' document member of the same patent family

2

Date of the actual completion of the international search

11 July 1996

Date of mailing of the international search report

25.07.96

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+31-70) 340-3016

Authorized officer

Perez Perez, J

INTERNATIONAL SEARCH REPORTInternational Application No
PCT/GB 96/00769**C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT**

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	INTERNATIONAL CONFERENCE ON COMMUNICATIONS, vol. 3, May 1993, CH, pages 1353-1357, XP000448363 R. KRISHNAN ET AL.: "An approach to path splitting in multipath networks" see paragraph 3 -----	1,10

2

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/GB 96/00769

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US-A-5291477	01-03-94	NONE	

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.